

SFA™ OS

Version 2.2.1.3

Product Release Notes

SFA7700 Series, SFA10K Series, and SFA12K Series

Important Information

Information in this document is subject to change without notice and does not represent a commitment on the part of DataDirect Networks, Inc. No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose other than the purchaser's personal use without the written permission of DataDirect Networks, Inc.

© 2014 DataDirect Networks, Inc. All rights reserved.

DataDirect Networks, the DataDirect Networks logo, DirectOS, DirectProtect, DirectMon, EXAScaler, GRIDScaler, Information in Motion, NAS Scaler, NoFS, ObjectAssure, SATAssure, Silicon Storage Appliance, S2A, Storage Fusion Architecture, SFA, Storage Fusion Fabric, Storage Fusion Xcelerator, SFX, xSTREAMScaler, Web Object Scaler, WOS are registered trademarks or trademarks of DataDirect Networks, Inc. All other brand and product names are trademarks of their respective holders.

DataDirect Networks makes no warranties, express or implied, including without limitation the implied warranties of merchantability and fitness for a particular purpose of any products or software. DataDirect Networks does not warrant, guarantee or make any representations regarding the use or the results of the use of any products or software in terms of correctness, accuracy, reliability, or otherwise. The entire risk as to the results and performance of the product and software are assumed by you. The exclusion of implied warranties is not permitted by some jurisdictions; this exclusion may not apply to you.

In no event will DataDirect Networks, their directors, officers, employees, or agents (collectively DataDirect Networks) be liable to you for any consequential, incidental, or indirect damages, including damages for loss of business profits, business interruption, loss of business information, and the like, arising out of the use or inability to use any DataDirect product or software even if DataDirect Networks has been advised of the possibility of such damages by you. Because some jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, these limitations may not apply to you. DataDirect Networks liability to you for actual damages from any cause whatsoever, and regardless of the form of the action (whether in contract, tort including negligence, product liability or otherwise), is limited to the sum you paid for the DataDirect product or software.

TABLE OF CONTENTS

1.0 OVERVIEW	4
2.0 INSTALLATION	4
2.1 Overview	4
2.2 Recommended and Required Code Levels	4
2.3 Procedure to Request Firmware	5
2.4 Pre-Firmware-Upgrade Procedure: SFA12K Series	5
2.5 Post-Firmware-Upgrade Procedure	7
SFA12K Series and SFA10K Series Products	7
3.0 ENHANCEMENTS	8
4.0 RESOLVED ISSUES	8
5.0 KNOWN ISSUES	8
5.1 Common	8
5.2 SFA12KXE and SFA12K-20E	10
5.3 SFA12KX, SFA12K-40, SFA12K-20, and SFA10K-X	12
5.4 SFA10K-X	14
5.5 SFA7700	14
5.6 Red Hat 6.2	15
6.0 RECOMMENDED PRACTICES	16
7.0 COMPATIBILITY GUIDE	17
8.0 SUPPORT	17

1.0 Overview

This document applies to the SFA OS version 2.2.1.3-21587 release for DataDirect Networks' SFA7700 Series, SFA10K Series, and SFA12K Series products. It details upgrade procedures, enhancements, resolved issues, known issues with workarounds, and recommended practices associated with SFA OS 2.2.1.3.

In this document, except where otherwise specified, SFA10K represents SFA10K, SFA10K-X, SFA10K-M, and SFA10K-E. SFA12K represents SFA12K-40, SFA12K-20, and SFA12K-20E. SFA12KX represents SFA12KX and SFA12KXE.

NOTE: SFA OS 2.2.0.2 changed the BBU (battery backup) lifespan from two to four years. Because battery lifespan cannot be accounted for while in inventory, beginning with SFA OS 2.2.1.3 on SFA12K Series and SFA10K series systems, the BBU “in service” life has been changed to three years. Customers should replace their BBU after three years of “in service” life.

2.0 Installation

2.1 Overview

SFA OS version 2.2.1.3 is a mandatory release for all supported systems.

NOTE: If you are upgrading from SFA OS v1.5.3, please see the “Upgrading 1.5.3” bullet in Section 5.1 “Common” before upgrading to SFA OS 2.x.

NOTE: The SFA OS 2.2 upgrade can take up to 25 minutes to complete. Additional time is required for any other firmware updates.

2.2 Recommended and Required Code Levels

The recommended and required controller code levels are shown below. Required code levels are marked with an asterisk (*).

SFA OS Products	BIOS	BMC	CONFIG	FPGA	SEP
SFA7700	24.001*	21.97*	n/a	002.003-000.049*	0110-000*
SFA10K	GEMDV250	DDGEN050	JANUSC23	n/a	n/a
SFA12K	21.0L	21.72 ¹	n/a	n/a	n/a
SFA12KX	21.0p*	21.72 ¹	n/a	n/a	n/a

NOTE: It is strongly recommended that SFA7700 systems are updated to the code levels shown above before installing SFA OS 2.2.x.

NOTE: Customers should not upgrade to BIOS 21.0p without contacting DDN Support first.

¹ This is automatically updated when SFA OS v2.2.0 is installed.

2.3 Procedure to Request Firmware

To request the new firmware, contact DDN by sending an email to support@ddn.com and include the following information:

Name:
Title:
Company:
Address:
Phone:
Email:
System Type:
System Serial Number:
Firmware Upgrade Requested:

A response will be sent to the provided email with instructions on downloading the requested firmware.

2.4 Pre-Firmware-Upgrade Procedure: SFA12K Series

NOTE: Before upgrading SFA OS on an SFA12K system, please read this section carefully. This release of firmware may update BMC and/or ICL firmware (if applicable), taking 10-35 minutes longer than a typical firmware update if either or both will be updated. The updates, if applicable, cannot be bypassed and must not be interrupted. An extra reboot (automatically performed by the firmware) may also occur to save and then load 'ui' configuration parameters (see Section 3.0 “Enhancements”). The remainder of this section shows how to verify if the updates will apply to your controllers.

SFA OS 2.2.1.3 has a feature that will upgrade the SFA12K Baseboard Management Controller (BMC) firmware version automatically as part of the SFA OS upgrade. The automatic update will only function with SFA12K BMC versions newer than 21.54.

Prior to upgrading to 2.2.1.3, please identify your current BMC version and then check these notes:

1. If you have BMC version 21.54 or older, you must perform a manual BMC firmware update before you load the SFA OS 2.2.1.3 firmware. Refer to the *SFA12K BIOS and BMC Firmware Update Field Bulletin*, 96-00344-001.
2. The automatic update of the BMC will take as long as 10 minutes. During that time, the SFA12K controller will be unresponsive and will not show progress. Including related steps, the overall time is approximately 25 minutes per controller.

WARNING: Do NOT interrupt this upgrade process; doing so could damage the controller.

3. If the BMC version is lower than 21.72, the BMC firmware will be automatically updated.

To identify your current BMC version:

1. View the output from '**SHOW ENCLOSURE**' to identify the enclosures representing your SFA12K chassis.

For example, see Figure 1.

Figure 1

```
$ show enclosure
```

```
*****
*      Enclosure(s)      *
*****
```

Idx	Pos	Type	Logical ID	Vendor ID	Product ID	Revision	...
0	NA	CONTROLLER	0x0001ff0900160000	DDN	SFA12000	0000	...
1	3	DISK	0x50001ff211320000	DDN	SS7000	0502	...
2	2	DISK	0x50001ff2114be000	DDN	SS7000	0502	...
3	5	DISK	0x50001ff2114ce000	DDN	SS7000	0502	...
4	1	DISK	0x50001ff21156a000	DDN	SS7000	0502	...
5	4	DISK	0x50001ff2114ba000	DDN	SS7000	0502	...
6	NA	CONTROLLER	0x0001ff0900100000	DDN	SFA12000	0000	...

```
Total Enclosures: 7
```

- For each enclosure Idx (Index) number showing 'Type' value of "Controller" (there should be two), issue '**SHOW ENCLOSURE <x> ALL**' to identify your current BMC version. For the output in Figure 1, you would use 'show enclosure 0 all' and 'show enclosure 6 all'. Example output for 'SHOW ENCLOSURE 0 ALL' is shown in Figure 2.

Figure 2

```
$ show enclosure 0 all
```

```
*****
*      Enclosure(s)      *
*****
```

Index:	0
Enclosure:	0x50000000
Type:	CONTROLLER
Logical ID:	0x0001ff0900160000
Vendor ID:	DDN
Product ID:	SFA12000
Revision:	0000
Is Responsive:	TRUE
Fault Status:	OK
Child Fault Status:	NON-CRITICAL
Zones:	0
Zoning Configuration:	0
Position:	NOT APPLICABLE
Fault Indicator:	OFF
Predicted Failure Ind:	OFF
Locate Indicator:	OFF
Baseboard part number:	X9DRX+-F-DD007
Baseboard serial number:	WM21S31110
Baseboard date:	19960101
BIOS version:	21.0k
BMC version:	21.67
Firmware version:	1.5.1.2

```
Total Enclosures: 1
```

- After the firmware is updated, the controller will reboot. As it performs the system check, it will check the ICL firmware version. If the ICL firmware version is older than the versions shown below, the SFA will automatically update the ICL of the local controller, and then it will reboot again.

- SFA12K Series: 2.31.5050
- SFA10K Series 2.9.1200

How to identify new vs. old firmware

The command **SHOW ICL_IOC** will report firmware version 2.31.5050 on updated systems.

New version:

```
SFA12K[0]$ show icl_ioc
*****
* ICL IOC(s) *
*****
|ICL|
Idx|Ctrlr|RP|IOC| Type |Vendor|Product| Slot | Model Name | HW Version | FW Version...
-----
0 1 0 00 IB HCA 15b3 1003 2 MT4099 0x00000000 2.30.5050...
```

Old version:

```
|ICL|
Idx|Ctrlr|RP|IOC| Type |Vendor|Product| Slot | Model Name | HW Version | FW Version...
-----
0 1 0 00 IB HCA 15b3 1003 2 MT4099 0x00000000 2.30.3000...
```

4. Finally, the InfiniBand cards have their firmware updated if they don't have the latest version. It takes less than a minute per card (so four cards would take less than four minutes) and one reboot.

2.5 Post-Firmware-Upgrade Procedure

SFA12K Series and SFA10K Series Products

After upgrading the firmware, perform these steps to initialize the battery life remaining feature:

1. Issue the command, **SHOW UPS * ALL_ATTRIBUTES**.
2. If you see a date in “Battery Mfg. Date”, do nothing more, this procedure is complete.
3. If you see the message shown in Figure 3, then issue the command, **SET UPS <id> BATTERY_MANUFACTURE_DATE=MM-DD-YYYY**, where “**id**” is the ID of the UPS that needs to be given a date, and **MM-DD-YYYY** is the date when the UPS battery was put into service (typically the install date of the subsystem).

Figure 3

```
Battery Mfg. Date: NOT AVAILABLE
Battery Life Remaining: NOT AVAILABLE
```

4. Issue the command, **SHOW UPS * ALL_ATTRIBUTES**. You should see the message shown in Figure 4.

Figure 4

```
Battery Mfg. Date: Thu Sep 8 4:10:30 2011
Battery Life Remaining: 730 days
```

3.0 Enhancements

SFA OS 2.2.1.3 has no enhancements.

4.0 Resolved Issues

The following issues are fixed in SFA OS v2.2.1.3.

- Tracking of preexisting bad blocks could be lost when upgrading to SFA OS 2.2.1.0 or 2.2.1.2.
- Previously, an error in the kernel memory allocator could cause a controller reboot on SFA10K-E, SFA12K-20E, and SFA12KXE systems. This happened rarely. [SR35724]
- CVE-2014-6271, CVE-2014-7169: bash code injection security vulnerability via environment variables (aka Shellshock).

The following issues are fixed in SFA OS v2.2.1.2.

- Resolved issues on embedded models related to application stacks (virtual machines) configured to Auto Start:
 - Following a controller boot-up, stacks could previously fail to auto start, while shown as “running.” (Could occur on any prior firmware release if Auto Start was enabled.)
 - Manual recovery or workaround prior to loading this firmware to address the issue: After waiting at least two minutes for VM boot, attempt to access the stack/VM console (method is configuration dependent; see the SFA OS User Guide). If the VM is found to not be accessible, **AP SHUTDOWN STACK <#> ABRUPT** may be used, followed by **AP START STACK <#>**. Note, improper use of the **AP SHUTDOWN** command (such as issuing this command on stacks which are on-line) may impact system operation or availability.
 - Following a controller boot or restart, an additional reboot could occur (rare).

5.0 Known Issues

5.1 Common

- **Upgrading 1.5.3:** When upgrading an SFA controller from 1.5.3 to 2.x, even though the upgrade process has completed, the second controller may not shut down and restart.

Once the upgrade process is started on the first controller, refrain from issuing any application commands with the **COUNTER** option from either controller until **after** the completion of the upgrade and restart of the second controller. Otherwise, the second controller may not shut down and restart with the new firmware.

NOTE: A controller will not attempt to shut down until it has completed the flushing of its mirrored cached data which may take up to an hour depending on I/O load and I/O transfer sizes.

- The drive serial number reported in **SHOW PD** output may contain additional characters that are not on the disk label. The serial number on the disk label will always be contained within the potentially more verbose string reported by SFA OS.
- The SS7000 enclosure does not support the **SET SLOT xx POWER OFF** command with enclosure firmware versions prior to 05.02.01.
- In the SS7000, when removing an I/O module or cable under load, it is possible that one or more disks could be marked failed which would change the redundancy of the storage pools. This is due to in-flight I/O being timed out and retried before the IOC reports the drive as missing. There is a high probability that the newly failed drives are fine and can be returned to service by clearing the failed state and reassigning them to pools.

Use the following procedure to clear a failed drive and reassign it back to a pool:

show unassigned failed	←Shows which PDs have failed; get drive's <id>
show pool	←Shows which pools are involved
clear pd <id> failed	←<id> is from show unassigned output
assign pd <id> to pool-id set	←Pool-ID is shown in the show pool output

NOTE: If the pool sparing policy is set to "SWAP", the SFA will typically auto rebuild the drives, using fractional or full rebuild depending on the duration of the interruption.

- If a drive is declared as Unknown or otherwise is missing from a pool and you choose the GUI to spare a drive into the pool, the drive size may display as needing a drive of 0 GB size.
- If stack commands (CLI commands starting with "app") are issued shortly after rebooting a controller, you may encounter communication errors displayed in the CLI, and failure of the attempted commands. Errors may include:
 - Long delays followed by:
Communication connection failed for this command
 - Or this message:
ASM initialization in progress

Suggested Work Around

- After the failed controller comes back up, do not immediately issue any CLI commands besides "show controller".
- Connect to the remaining controller and issue this command:
\$ show controller
- When "**SHOW CONTROLLER**" shows the remote controller in the output, focus on the ULA field in the output. The remote controller will not show up until it is fully booted.

- If the ULA field has the value "0000000000000000", continue to wait and not issue any CLI commands besides "**SHOW CONTROLLER**".

Below is an example output of the condition where you would want to wait:

```
$ show controller

*****
*      Controller(s)      *
*****

      |      Up Time      |      |Encl|      |...
Idx|Name      |Mastership|Locality| D: H: M: S|RP|      ID      |Idx |      ULA      |...
-----|-----|-----|-----|---|-----|-----|-----|-----|
  0 A      PRIMARY    LOCAL   0000:20:24:34  1 0001ff0900180000  0 00000001ff0800ac ...
  1 B      SECONDARY  REMOTE  0000:00:00:05  1 0001ff09002d0000  0 0000000000000000 ...

Total Controllers: 2
```

- When the ULA number has something other than all zeros for the remote controller, it is now safe to issue CLI commands and avoid the communication problems.
- There is an issue when enabling DIRECTPROTECT of DATA_INTEGRITY_FIELD for a pool that is already doing an initialization.

When DATA_INTEGRITY_FIELD is enabled after an initialization has already started, that initialization is stopped and restarted from the beginning as a Non-Destructive Initialization to fill in the DATA_INTEGRITY_FIELD blocks. If, for any reason other than a shutdown, the master controller were to restart during the transition from a DIRECTPROTECT of non-DATA_INTEGRITY_FIELD to DATA_INTEGRITY_FIELD the Initialization fence may not get set back to zero. If this were to happen the Non-Destructive Initialization would not properly fill out the DATA_INTEGRITY_FIELD blocks below the previous Initialization fence. This is a very small window while the controller is in the process of transitioning to the new DIRECTPROTECT value.

Work Arounds:

- When creating the pool, specify the DIRECTPROTECT value on the command line itself. **CREATE POOL [options needed] DIRECTPROTECT=DATA_INTEGRITY_FIELD.**
- Wait for the original initialization to complete prior to setting DIRECTPROTECT to DATA_INTEGRITY_FIELD.

5.2 SFA12KXE and SFA12K-20E

- On E-platforms, downgrading SFA OS 2.2.0 to earlier SFA OS versions is not supported if there are more than 128 VDs presented to any stack.
- If Installing Older DDN Solutions on SFA12K-20E, Potential Network Ordering Issues May Arise.

If you install a version of GRIDScaler older than 2.0.0 or EXAScaler older than 1.6.1 and you are using enhanced network management in the VMs (sr-iov virtual function interfaces), you could run in to an Ethernet ordering issue during installation where eth0 may not be associated with the expected driver. This could result in failure of user mode networking (the port 555x access) or other problems.

A workaround is to temporarily un-bind the virtual function interfaces from the stack using CLUI. Once installation is complete, re-associate the virtual function interfaces with the VMs and then configure networking.

IOC index 01792 is associated with the first SFA controller's first Ethernet port. IOC index 01793 is associated with the first SFA controller's second Ethernet port. IOC index 34560 is associated with the second SFA controller's first Ethernet port. IOC index 34561 is associated with the second SFA controller's second Ethernet port.

The following CLUI example would unbind the virtual function interfaces from stacks configured by default on systems leaving DDN:

```
app unbind stack 0 ioc 01792
app unbind stack 0 ioc 01793
app unbind stack 1 ioc 01792
app unbind stack 1 ioc 01793
app unbind stack 2 ioc 01792
app unbind stack 2 ioc 01793
app unbind stack 3 ioc 01792
app unbind stack 3 ioc 01793
app unbind stack 32768 ioc 34560
app unbind stack 32768 ioc 34561
app unbind stack 32769 ioc 34560
app unbind stack 32769 ioc 34561
app unbind stack 32770 ioc 34560
app unbind stack 32770 ioc 34561
app unbind stack 32771 ioc 34560
app unbind stack 32771 ioc 34561
```

At this point, to complete the installation, start up your stacks as you normally would using CLUI.

After installation is complete on all the VMs, you can re-bind the virtual function interfaces and shutdown/startup the VMs to make use of the interfaces. Here is an example using CLUI:

```
app bind stack 0 ioc 01792
app bind stack 0 ioc 01793
app bind stack 1 ioc 01792
app bind stack 1 ioc 01793
app bind stack 2 ioc 01792
app bind stack 2 ioc 01793
app bind stack 3 ioc 01792
app bind stack 3 ioc 01793
app bind stack 32768 ioc 34560
app bind stack 32768 ioc 34561
app bind stack 32769 ioc 34560
app bind stack 32769 ioc 34561
```

```
app bind stack 32770 ioc 34560
app bind stack 32770 ioc 34561
app bind stack 32771 ioc 34560
app bind stack 32771 ioc 34561
```

Now, shut down and start up the VMs/stacks again using CLUI. After they boot up, you can configure networking on the virtual machines.

- Suggested Ethernet Ordering in Virtual Machines for SFA12K-20E

For new installations, DDN recommends the following virtual machine Ethernet ordering starting in SFA OS 2.0.0.

- VM eth0: virtual function interface associated with controller's first port
- VM eth1: virtual function interface associated with controller's 2nd port
- VM eth2: user mode networking support (provides port 555x ssh functionality)

Ethernet interface numbers larger than eth2 may be associated with Mellanox cards operating in 10G/40G mode.

Controller eth0 (and VM eth0) are meant to be connected to the customer network and may also be used by DDN Solutions like GRIDScaler for network management. Controller eth1 (and VM eth1) are reserved for DirectMon.

Starting with EXAScaler 1.6.1 and GRIDScaler 2.0.0, this naming and usage convention is used and configured as part of the ISO installation process.

The *SFA OS User Guide* has more information on Ethernet interface configuration for controllers and VMs.

- On the SFA12K-20E only, ConnectX3 HCA IB read and write performance drops on PCI bus 1 when the message size is greater than 128 KiB.

5.3 SFA12KX, SFA12K-40, SFA12K-20, and SFA10K-X

- Some sequential read workloads with request sizes between 128 KB and 1024 KB may require enabling read ahead caching (prefetch) to obtain full performance. (`SET POOL <#|*> READ_AHEAD_CACHING ON`).

- To improve failover time with RHEL 6.x and Qlogic cards, the following is recommended to enable the QLogic driver to react quickly to a loss of path:

Qlogic driver: 8.04.00.06.0-k

Add the parameter 'ql2xasynctmfenable=1' to /etc/modprobe.d/modprobe.conf

Example line in modprobe.conf:

options qla2xxx ql2xfwloadbin=1 ql2xqfullrampup=15 ql2xasynctmfenable=1

- To improve failover time and to prevent I/O errors, the following settings in multipath.conf are recommended:

(These settings can be changed for the defaults section which will apply to all devices, or just under the SFA devices):

```
checker_timeout    5
dev_loss_tmo      10
fast_io_fail_tmo   5
```

These settings are included in the DDN multipath package version 1.5-5 and above.

- When the controller is preparing to shutdown, it will first put all its pools into write-through mode and attempt to flush all the dirty cache. On SFA platforms with multiple RAID processors (RP), in the case where one RP finishes flushing its cache before the other RP, the first RP to finish flushing its cache will not service I/O from the host until the SFA reboots. This may cause I/O errors on the host and cause applications on the hosts to hang and eventually time out. To work around this issue, reduce I/O load during planned maintenance activities such as firmware upgrades and reboots of the controllers so that the flush activities will complete quickly.
- With RHEL6.2 and OFED 1.5.4.0 in an IB switch attached environment, it is possible that a virtual disk on a controller may not be added back to the multipath device map after a failover.

To find the offline device, issue the command:

```
lsscsi | awk -F/ '{print $NF}' | while read a; do printf "%s\n" $a; cat /sys/block/$a/device/state; done
```

To bring the device back online, issue the command:

```
#echo running > /sys/block/<sd??>/device/state and #multipath -r
```

where you replace <sd??> with the appropriate sd, for example *sdaf*, found from the previous command.

A workaround for this issue is to update these packages:

- **device-mapper:1.02.74-10.e16**
- **device-mapper-multipath: 0.4.9-56.el6_3.1.x86_64**

DDN recommends that you install these as soon as possible.

- In an InfiniBand switch environment, there is a small chance that if a cable between the switch and the controller is pulled, both the physical link and the logical link will be lost. This has occurred with a Mellanox 6025F Switch and a Mellanox HCA.
- In an InfiniBand switch environment running RHEL 5.7, if a cable is pulled from either an initiator or a target, a failover occurs as expected; however, once the connection is reestablished, it does not fail back to the original controller.
 - In order to resolve this issue, you must issue the following command from the initiator:
udevtrigger.
 - The Linux man pages state that the command will simply “request kernel devices events for coldplug”. This will force udev to send a notification for newly discovered path, which will allow the multipath daemon to detect that the path has returned. After the multipath discovery takes place, I/Os can be rebalanced back to allow the preferred paths to be used.
- In an InfiniBand switch environment, an online upgrade causes an I/O error while upgrading the second controller with SLES 10 SP4. In an FC environment, there are no issues because the discovery rate is much quicker. The only option for SFA upgrades with SLES 10 SP4 is to perform the upgrade offline.

To perform the upgrade offline,

1. Issue the CLI command **SET SUBSYSTEM OFFLINE.**

2. Perform the controller firmware upgrade on both controllers.
3. Issue the CLI command **SHUTDOWN SUBSYSTEM RESTART**.
4. After both controllers come back online, issue the CLI command **APP SHOW DISCOVERED_INITIATOR**. This should report “No discovered_initiators exist.” This will be the state until a read request for information from those SCSI devices is issued. Once the devices are accessed, there will be a small delay as they are rediscovered by the system.

5.4 SFA10K-X

- The SS2460 does not support SLOT POWER OFF feature and so the **SET SLOT POWER** command has no effect on drives in the SS2460.
- The Management Ethernet port (LAN4) may boot running at 10 Mb/sec half-duplex, instead of 100 Mb/sec or 1000 Mb/sec full-duplex. Typically, full Ethernet port speeds can be restored by unplugging and replugging the Ethernet cable to the controller port experiencing the slow speed. If the issue persists, please contact, support@ddn.com.
- On SS6000 enclosures the DISPLAY POWER_SUPPLY command may fail to display the last three characters of the power supply serial number. This is only a display issue, the command will complete successfully with no adverse effects other than the truncated results.
- The SS6000 may encounter a problem where show enclosure commands will return the following error message: “A command issued to an enclosure to get attributes for the UI timed out.” This is an issue in the enclosure firmware v3.001 that is resolved by upgrading to enclosure firmware v3.010.
- Powering off drives is not supported with 6KSSI interposers in the SS6000 enclosure. If you issue the command, **SET SLOT <encl-id slot-id> POWER OFF**, the system erroneously displays a success message. However, the drive power is not affected by the command.

5.5 SFA7700

- **After a controller shutdown the BMC WebUI indicates that the power is on.** If you are using the BMC WebUI remote power feature, then after a SHUTDOWN, CONTROLLER, or SHUTDOWN SUBSYSTEM the WebUI will indicate that the Host power is still on and the button for turning the power back on is grayed out.

Resolution: Switch to another page in the WebUI and then back to the remote power control page, this will update the page correctly.

- **There is no way to enable the BMC IP port once it is disabled with the WebUI.** It is possible to turn off the BMC IP port with the WebUI, but once it is off it cannot be turned on by the CLUI or ipmitool.

Resolution: Don't disable the BMC IP port. Contact DDN support if you do.

- **IPMI Power Off does not always turn the canister completely off.**

After a **SHUTDOWN CONTROLLER** or **SHUTDOWN SUBSYSTEM**, which ends with an IPMI Power Off, the canister fan may restart even though the canister is in the shutdown or standby state.

Resolution: This condition is benign except that it makes it confusing whether the canister is shutdown or not. If all five canister LEDs are off, then it did indeed shut down (trust the LEDs).

- Occasionally an extra enclosure shows up in **SHOW ENCLOSURE** immediately after powering up the system.

Resolution: Shut the system down, power-cycle, and start it up again.

- With SFA7700 BMC/FPGA 21.97/2.3-0.49 the reported battery charge will not rise to 100%. Maximum charge levels may be as low as 89%. This should not affect the ability of the battery pack to hold up the canister during a power failure. The issue will be fixed in the next BMC/FPGA release.

5.6 Red Hat 6.2

This section applies to the SFA12K-40, SFA12K-20, SFA10K-X, and SFA10K-M.

- Updating the kernel and multipath packages to the latest version is recommended. As of this release, minimum DDN qualified versions include:
 - the kernel version - 2.6.32-279.2.1.el6.x86_64
 - the multipath version - 0.4.9-56.el6_3.1.x86_64.
- The kernel modules that come with OFED 1.5.4 will cause I/O threads to hang when the data path between the controller and switch is lost. A workaround for this issue is to use the Red Hat supplied InfiniBand solution. For a switched InfiniBand environment:
 - Install the in-box OFED drivers that come with RHEL 6 Update 2
 - Update the multipath-tools package to a minimal version of 0.4.9-56.el_3.1
 - Do not use user-friendly names in multipath.conf
 - Configure RHEL 6 Update 2 to a minimal kernel version of 2.6.32-279.2.1.el6
- These packages are required for InfiniBand support on Red Hat. (Use the command **yum groupinstall infiniband** to install all IB packages):
 - ibutils-1.5.7-6.el6.x86_64
 - infiniband-diags-1.5.12-3.el6.x86_64
 - opensm-libs-3.3.13-1.el6.x86_64
 - opensm-3.3.13-1.el6.x86_64
 - libmlx4-1.0.2-5.el6.x86_64
 - srptools-0.0.4-15.el6.x86_64
- To start IB lun discovery, issue these commands:
 1. **modprobe ib_umad**

2. **modprobe ib_uverbs**
3. **modprobe mlx4_ib**
4. **opensm -g <port GUID>** (where *port GUID* is the port GUID from ibstat output)
5. **run_srp_daemon -R 20 -T 10 -nce -i <ca_name> -p <port>** (where *ca_name* is the channel adapter name and *port* is the port number)

6.0 Recommended Practices

The recommended best practices for SFA OS version 2.2.1.3 include:

- **SS6000 DEM and I/O module replacement.**

For SS6000 DEM and I/O module replacement, see the SFA OS User Guide for the procedures.

- **On hosts that are directly connected, if the configuration is cleared, the Fibre Channel initiators may not log back into the subsystem.**

This results in no initiators being available to assign to new presentations. To correct this, do a rescan from the host before setting up your presentations.

This does not apply to the SFA12K-40 and SFA12K-20.

- **Disk firmware upgrades.**

Before performing a disk firmware upgrade, scp the .DDN file to both controllers.

This avoids a situation where a controller attempts to perform the upgrade but can't find the .DDN file.

- **Firmware upgrades should be done during low usage periods whenever possible.**

As upgrading your system will interrupt I/O operations as paths fail over, be aware of the possibility of increased I/O latencies during the upgrade operation. Prior to upgrading your firmware, DDN recommends reviewing the system logs to ensure the system is running cleanly. DDN also recommends allowing any rebuilds and forced verify jobs to complete and pausing any running verify jobs.

When a controller is shut down, all cached data must be flushed before the controller will shut down. This operation can take a considerable amount of time depending upon the amount of cache and if the cached data is sequential vs. random, large block vs. small block, number of VDs, and so on. Should the shutdown fail, an event will be generated. If you do not see the event, the controller is still shutting down. Please allow it to complete to avoid possible data loss issues.

- **Update the secondary controller first.**

When performing an SFA OS upgrade, upgrade the secondary controller first. Doing this requires primary to move once instead of twice.

7.0 Compatibility Guide

The *SFA OS Compatibility Guide* lists the client operating systems, HBA and HCAs, enclosures, and disk drives that are supported by SFA OS. The guide is available at <http://www.ddn.com/en/support/product-downloads-and-documentation>.

NOTE: The Compatibility Guide applies to SFA12K Series, SFA10K Series, and SFA7700 Series.

8.0 Support

To submit debug information to DDN, use the **diag tgz** command.

Diag is a captive account on SFA OS systems that is used for performing operations on a controller when there's no CLI access or the CLI isn't appropriate. Perform diag operations by using ssh to login to the diag account.

The command syntax is:

```
ssh diag@xxx.xxx.xxx.xxx diag-command
```

(where xxx.xxx.xxx.xxx is the controller's IP address and diag-commands are listed in Table 1.)

Table 1 Valid diag commands and descriptions

Command	Description
help	Displays the help information.
clear	Clears all old crash files (it will clear a recursive crash).
reboot	Used to restart the controller if the RAID stack is not running. If the RAID stack is running, this command will fail and you should use the CLI to perform this operation.
shutdown	Used to shutdown the controller if RAID stack is not running. If the RAID stack is running, this command will fail and you should use the CLI to perform this operation.
tgz	Creates a .tgz file of all the crash information. The file is output to stdout and since it's a binary file, having it display on your terminal isn't desirable. Redirect stdout to your workstation with a command like: ssh diag@xxx.xxx.xxx.xxx tgz > crash.tgz

Please contact DataDirect Networks Support at any time for assistance. Support can be reached by these methods:

Web

Support Portal <https://portal.ddnsupport.com>
Portal Assistance webportal.support@ddn.com
Technical Support Bulletins <http://www.ddn.com/en/technical-support-bulletins>

Email

Support Email support@ddn.com
Mailing List Subscriptions support-tsb@ddn.com

Telephone

DDN Worldwide Directory <http://www.ddn.com/en/support/contact-support>

NOTE: If you are a beta test site, when contacting DDN, identify yourself as a beta test site so DDN can dispatch your request accordingly.
